

INTRODUCTION TO POMP BY EXAMPLE

AARON A. KING

1. A FIRST EXAMPLE: A TWO-DIMENSIONAL RANDOM-WALK.

In order to specify a partially-observed Markov process, we must define the process model and the measurement model. In particular, we will need to be able to simulate from and compute the p.d.f. of both models. The following function will simulate the process model. The documentation (`?pomp`) spells out the specifications for this function.

```
> rw.rprocess <- function(xstart, times, params,
+   ...) {
+   nsims <- ncol(params)
+   ntimes <- length(times)
+   dt <- diff(times)
+   x <- array(0, dim = c(2, nsims, ntimes))
+   rownames(x) <- rownames(xstart)
+   noise.sds <- params[c("s1", "s2"), ]
+   x[, , 1] <- xstart
+   for (j in 2:ntimes) {
+     x[, , j] <- rnorm(n = 2 * nsims, mean = x[,
+       , j - 1], sd = noise.sds * dt[j -
+       1])
+   }
+   x
+ }
```

Some methods will require the probability density of a given state transition. The function `dprocess` will evaluate this for a sequences of state transitions.

```
> rw.dprocess <- function(x, times, params, log = FALSE,
+   ...) {
+   nsims <- ncol(params)
+   ntimes <- length(times)
+   dt <- diff(times)
+   d <- array(0, dim = c(2, nsims, ntimes - 1))
+   noise.sds <- params[c("s1", "s2"), ]
+   for (j in 2:ntimes) d[, , j - 1] <- dnorm(x[,
+     , j] - x[, , j - 1], mean = 0, sd = noise.sds *
+     dt[j - 1], log = TRUE)
+   if (log) {
+     apply(d, c(2, 3), sum)
+   }
+   else {
+     exp(apply(d, c(2, 3), sum))
+   }
+ }
```

Now we specify the function that will simulate the measurement process. Again, the documentation spells out how.

```
> bvnorm.rmeasure <- function(x, times, params,
+   ...) {
+   nsims <- dim(x)[2]
+   ntimes <- dim(x)[3]
+   y <- array(0, dim = c(2, nsims, ntimes))
+   rownames(y) <- c("y1", "y2")
+   for (k in 1:nsims) {
+     for (j in 1:ntimes) {
+       y[, k, j] <- rnorm(2, mean = x[, k,
+         j], sd = params["tau", k])
+     }
+   }
+   y
+ }
```

Finally, we have to specify how to evaluate the likelihood of an observation given the underlying state.

```
> bvnorm.dmeasure <- function(y, x, times, params,
+   log = FALSE, ...) {
+   d1 <- dnorm(x = y["y1", ], mean = x["x1",
+     ], sd = params["tau", ], log = TRUE)
+   d2 <- dnorm(x = y["y2", ], mean = x["x2",
+     ], sd = params["tau", ], log = TRUE)
+   if (log) {
+     sum(d1, d2, na.rm = T)
+   }
+   else {
+     exp(sum(d1, d2, na.rm = T))
+   }
+ }
```

The following builds a pomp object called `rw2`.

```
> rw2 <- pomp(rprocess = rw.rprocess, dprocess = rw.dprocess,
+   rmeasure = bvnorm.rmeasure, dmeasure = bvnorm.dmeasure,
+   times = 1:100, data = rbind(y1 = rep(0, 100),
+     y2 = rep(0, 100)), t0 = 0, useless = 23)
```

Now we'll specify some parameters and initial states.

```
> p <- rbind(s1 = c(2, 2, 3), s2 = c(0.1, 1, 2),
+   tau = c(1, 5, 0))
      [,1] [,2] [,3]
s1    2.0   2   3
s2    0.1   1   2
tau   1.0   5   0

> x0 <- rbind(x1 = c(0, 0, 5), x2 = c(0, 0, 0))
      [,1] [,2] [,3]
x1     0   0   5
x2     0   0   0
```

Each column is a different initial state or parameter vector. Note that we must use rownames!

When we defined `rw2`, the data were all missing. We can generate simulated data by:

```
> examples <- simulate(rw2, xstart = x0, params = p)
> rw2 <- examples[[1]]
```

By default `simulate` will generate a list of new `pomp` objects. It can also be used to obtain the simulated state and/or measurement trajectories:

```
> y <- simulate(rw2, xstart = x0, params = p, obs = T,
+             states = T)
> y <- simulate(rw2, xstart = x0, params = p, obs = T)
> x <- simulate(rw2, xstart = x0, params = p, states = T)
> x <- simulate(rw2, nsim = 10, xstart = x0, params = p,
+             states = T)
> x <- simulate(rw2, nsim = 10, xstart = x0[, 1],
+             params = p[, 1], states = T)
> x <- simulate(rw2, nsim = 10, xstart = x0[, 1],
+             params = p[, 1], obs = T, states = T)
> x <- simulate(rw2, nsim = 10, xstart = x0, params = p[,
+             1], obs = T, states = T)
> x <- simulate(rw2, nsim = 10, xstart = x0[, 2],
+             params = p, obs = T, states = T)
> x <- simulate(rw2, nsim = 10, xstart = x0[, 1],
+             params = p[, 1])
```

A plot method exists for `pomp` objects (Fig. 1).

Access to the individual components of the `pomp` object is available by means of a few *methods*. To extract the data and the observation times, use `data.array` and `time`, respectively:

```
> x <- data.array(rw2)
> t <- time(rw2)
```

To run the process model, users should use `simulate` with the `states=T` option. Lower-level access is available via the `rprocess` method:

```
> x <- rprocess(rw2, xstart = x0, times = 0:100,
+             params = p)
```

Similarly, low-level access to the measurement-model simulator can be had through `rmeasure`:

```
> y <- rmeasure(rw2, x = x[, , -1, drop = F], times = 1:100,
+             params = p)
```

Access to the process-model p.d.f. is available via `dprocess`:

```
> dprocess(rw2, x[, , 6:11], times = 5:10, params = p)
```

```
      [,1]      [,2]      [,3]      [,4]
[1,] 0.44706134 0.39210511 0.603803022 0.45406540
[2,] 0.06985799 0.07635331 0.039727353 0.01215335
[3,] 0.02058891 0.01200745 0.005803537 0.01486791
      [,5]
[1,] 0.77678661
[2,] 0.07143071
[3,] 0.01871387

> dprocess(rw2, x[, , 6:11], times = 5:10, params = p,
+         log = T)
      [,1]      [,2]      [,3]      [,4]      [,5]
[1,] -0.8050595 -0.9362253 -0.5045073 -0.789514 -0.2525896
[2,] -2.6612908 -2.5723839 -3.2257153 -4.410150 -2.6390274
[3,] -3.8830025 -4.4222280 -5.1492877 -4.208550 -3.9784906
```

Note that `dprocess` returns a matrix: the rows correspond to independent simulations, the columns to distinct state transitions. The measurement-model p.d.f. is accessed via the `dmeasure` method, which like `dprocess`, returns a matrix. The rows correspond to independent simulations, the columns to distinct times.

```
> plot(rw2)
```




FIGURE 1. A plot method exists for `pomp` objects.

```
> dmeasure(rw2, y = y[, 1, 1:4], x = x[, , 2:5,
+         drop = F], times = time(rw2)[1:4], p)
[1] 0
> dmeasure(rw2, y = y[, 2, 1:4], x = x[, , 2:5,
+         drop = F], times = time(rw2)[1:4], p)
[1] 0
> dmeasure(rw2, y = y[, 3, 1:4], x = x[, , 2:5,
+         drop = F], times = time(rw2)[1:4], p)
[1] NaN
> exp(dmeasure(rw2, y = y[, 3, 1:4], x = x[, , 2:5,
+         drop = F], times = time(rw2)[1:4], p, log = T))
[1] NaN
```

2. A TWO-DIMENSIONAL ORNSTEIN-UHLENBECK PROCESS.

To keep things simple, we will study a discrete-time process. The tricks below will continue to be useful even in the case of a continuous-time process, but the computational effort will be greater. The unobserved Ornstein-Uhlenbeck (OU) process $X_t \in \mathbb{R}^2$ satisfies

$$X_t = A X_{t-1} + \xi_t.$$

The observation process is

$$Y_t = B X_t + \varepsilon_t.$$

In these equations, A and B are 2×2 constant matrices; ξ_t and ε_t are mutually-independent families of i.i.d. bivariate normal random variables. We let $\sigma\sigma^T$ be the variance-covariance matrix of ξ_t , where σ is lower-triangular; likewise, we let $\tau\tau^T$ be that of ε_t .

We build the `pomp` object by specifying the three basic elements. The process model simulator and density functions:

```
> ou2.rprocess <- function(xstart, times, params,
+   ...) {
+   nsims <- ncol(xstart)
+   ntimes <- length(times)
+   alpha <- array(params[c("alpha.1", "alpha.2",
+     "alpha.3", "alpha.4"), ], dim = c(2, 2,
+     nsims))
+   sigma <- array(params[c("sigma.1", "sigma.2",
+     "sigma.2", "sigma.3"), ], dim = c(2, 2,
+     nsims))
+   sigma[1, 2, ] <- 0
+   x <- array(0, dim = c(2, nsims, ntimes))
+   rownames(x) <- rownames(xstart)
+   x[, , 1] <- xstart
+   for (k in 1:nsims) {
+     for (j in 2:ntimes) {
+       x[, k, j] <- alpha[, , k] %*% x[,
+         k, j - 1] + sigma[, , k] %*% rnorm(2)
+     }
+   }
+   x
+ }
```

```
> ou2.dprocess <- function(x, times, params, log = FALSE,
+   ...) {
+   nsims <- ncol(x)
+   ntimes <- length(times)
+   alpha <- array(params[c("alpha.1", "alpha.2",
+     "alpha.3", "alpha.4"), ], dim = c(2, 2,
+     nsims))
+   sigma <- array(params[c("sigma.1", "sigma.2",
+     "sigma.2", "sigma.3"), ], dim = c(2, 2,
+     nsims))
+   sigma[1, 2, ] <- 0
+   d <- array(0, dim = c(nsims, ntimes - 1))
+   for (k in 1:nsims) {
+     for (j in 2:ntimes) {
+       z <- forwardsolve(sigma[, , k], x[,
+         k, j] - alpha[, , k] %*% x[, k,
+         j - 1])
+     }
+   }
```

```

+         if (log) {
+             d[k, j - 1] <- sum(dnorm(z, mean = 0,
+                                     sd = 1, log = TRUE))
+         }
+         else {
+             d[k, j - 1] <- exp(sum(dnorm(z,
+                                     mean = 0, sd = 1, log = TRUE)))
+         }
+     }
+ }
+ d
+ }

```

The measurement model is the same as that for the random walk example above. We build the `pomp` object:

```

> ou2 <- pomp(times = seq(1, 100), data = rbind(y1 = rep(0,
+ 100), y2 = rep(0, 100)), t0 = 0, rprocess = ou2.rprocess,
+ dprocess = ou2.dprocess, rmeasure = bvnorm.rmeasure,
+ dmeasure = bvnorm.dmeasure)

```

Now we'll specify the "true" parameters and initial states.

```

> x0 <- c(x1 = 50, x2 = -50)
x1 x2
50 -50
> p <- c(alpha.1 = 0.9, alpha.2 = 0, alpha.3 = 0,
+ alpha.4 = 0.99, sigma.1 = 1, sigma.2 = 0,
+ sigma.3 = 2, tau = 1)
alpha.1 alpha.2 alpha.3 alpha.4 sigma.1 sigma.2 sigma.3
0.90 0.00 0.00 0.99 1.00 0.00 2.00
tau
1.00

```

As before, we'll fill in the missing values with simulated data.

```

> tic <- Sys.time()
> ou2 <- simulate(ou2, xstart = x0, params = p,
+ nsim = 1000)[[1]]
> toc <- Sys.time()
> print(toc - tic)

```

Time difference of 8.123347 secs

Let's make sure everything works.

```

> x <- rprocess(ou2, xstart = as.matrix(x0), times = c(0,
+ time(ou2)), params = as.matrix(p))
> y <- rmeasure(ou2, x = x[, , -1, drop = F], times = time(ou2),
+ params = as.matrix(p))
> dprocess(ou2, x[, , 36:41, drop = F], times = time(ou2)[35:40],
+ params = as.matrix(p))
      [,1]      [,2]      [,3]      [,4]      [,5]
[1,] 0.1510728 0.0568224 0.07681458 0.00693994 0.1514396
> exp(dprocess(ou2, x[, , 36:41, drop = F], times = time(ou2)[35:40],
+ params = as.matrix(p), log = T))
      [,1]      [,2]      [,3]      [,4]      [,5]
[1,] 0.1510728 0.0568224 0.07681458 0.00693994 0.1514396
> dmeasure(ou2, y = y[, 1, 1:4], x = x[, , 2:5,
+ drop = F], times = time(ou2)[1:4], params = as.matrix(p))

```

```
[1] 3.04314e-05
> exp(dmeasure(ou2, y = y[, 1, 1:4], x = x[, , 2:5,
+   drop = F], times = time(ou2)[1:4], params = as.matrix(p),
+   log = T))
[1] 3.04314e-05
```

3. LINKING IN COMPILED CODE FOR COMPUTATIONAL EFFICIENCY.

The functions we've just written are (relatively) easy to read, but they will be slow to evaluate because, being written in R, they must be interpreted. Since many of the methods we will use require us to simulate the process and/or measurement models many times, it is a good idea to translate these codes into a compiled language. The package includes some C codes that were written to implement the OU example. Read the source (file 'ou2.c') for details. The following wrapper functions make use of these compiled codes.

```
> ou2.rprocess <- function(xstart, times, params,
+   ...) .Call("ou2_simulator", xstart, times,
+   params)
> ou2.dprocess <- function(x, times, params, log = FALSE,
+   ...) .Call("ou2_density", x, as.numeric(times),
+   params, log)
> bvnorm.dmeasure <- function(y, x, times, params,
+   log = FALSE, ...) .Call("bivariate_normal_dmeasure",
+   y, x, as.numeric(times), params, log)
> bvnorm.rmeasure <- function(x, times, params,
+   ...) .Call("bivariate_normal_rmeasure", x,
+   as.numeric(times), params)
```

To take advantage of the compiled functions, we need to reconstruct the `pomp` object.

```
> ou2 <- pomp(times = seq(1, 100), data = rbind(y1 = rep(0,
+   100), y2 = rep(0, 100)), t0 = 0, rprocess = ou2.rprocess,
+   dprocess = ou2.dprocess, rmeasure = bvnorm.rmeasure,
+   dmeasure = bvnorm.dmeasure, ivpnames = c("x1.0",
+   "x2.0"), parnames = c("alpha.1", "alpha.2",
+   "alpha.3", "alpha.4", "sigma.1", "sigma.2",
+   "sigma.3", "tau"))
```

The `pomp` object we just created is included in the package: use `data(ou2)` to retrieve it. We'll fill the data slot with simulated data:

```
> tic <- Sys.time()
> ou2 <- simulate(ou2, xstart = x0, params = p,
+   nsim = 1000)[[1]]
> toc <- Sys.time()
> print(toc - tic)
```

Time difference of 0.860604 secs

Notice that we have added two objects, `ivpnames` and `parnames` to the `pomp` object. These will be passed to each of the functions and will come in handy later when we do particle filtering. Fig. 2 plots the data.

Let's make sure everything works.

```
> x <- rprocess(ou2, xstart = as.matrix(x0), times = c(0,
+   time(ou2)), params = as.matrix(p))
> y <- rmeasure(ou2, x = x[, , -1, drop = F], times = time(ou2),
+   params = as.matrix(p))
> log(dprocess(ou2, x[, , 36:41, drop = F], times = time(ou2)[35:40],
+   params = as.matrix(p)))
```

```
> plot(ou2)
```




FIGURE 2. The OU process.

```

      [,1]      [,2]      [,3]      [,4]      [,5]
[1,] -2.711203 -1.941398 -2.775977 -2.037255 -4.755495
> dprocess(ou2, x[, , 36:41, drop = F], times = time(ou2)[35:40],
+   params = as.matrix(p), log = T)
      [,1]      [,2]      [,3]      [,4]      [,5]
[1,] -2.711203 -1.941398 -2.775977 -2.037255 -4.755495
> log(dmeasure(ou2, y = y[, 1, 1:4], x = x[, , 2:5,
+   drop = F], times = time(ou2)[1:4], params = as.matrix(p)))
      [,1]      [,2]      [,3]      [,4]
[1,] -2.226756 -1.849295 -1.898595 -2.277097
> dmeasure(ou2, y = y[, 1, 1:4], x = x[, , 2:5,
+   drop = F], times = time(ou2)[1:4], params = as.matrix(p),
+   log = T)
      [,1]      [,2]      [,3]      [,4]
[1,] -2.226756 -1.849295 -1.898595 -2.277097

```


4. PARTICLE FILTER.

We can run a particle filter as follows:

```
> X0 <- matrix(x0, 2, 2000)
> rownames(X0) <- c("x1", "x2")
> fit1 <- pfilter(ou2, X0, p, filter.mean = T, pred.mean = T,
+   pred.var = T)
```

We can compare the results against those of the Kalman filter, which is exact in this case. First, we need to implement the Kalman filter.

```
> kalman.filter <- function(y, x0, a, b, sigma,
+   tau) {
+   n <- nrow(y)
+   ntimes <- ncol(y)
+   sigma.sq <- sigma %% t(sigma)
+   tau.sq <- tau %% t(tau)
+   inv.tau.sq <- solve(tau.sq)
+   cond.dev <- numeric(ntimes)
+   filter.mean <- matrix(0, n, ntimes)
+   pred.mean <- matrix(0, n, ntimes)
+   pred.var <- array(0, dim = c(n, n, ntimes))
+   dev <- 0
+   m <- x0
+   v <- diag(0, n)
+   for (k in seq(length = ntimes)) {
+     pred.mean[, k] <- M <- a %% m
+     pred.var[, , k] <- V <- a %% v %% t(a) +
+       sigma.sq
+     q <- b %% V %% t(b) + tau.sq
+     r <- y[, k] - b %% M
+     cond.dev[k] <- n * log(2 * pi) + log(det(q)) +
+       t(r) %% solve(q, r)
+     dev <- dev + cond.dev[k]
+     q <- t(b) %% inv.tau.sq %% b + solve(V)
+     v <- solve(q)
+     filter.mean[, k] <- m <- v %% (t(b) %%
+       inv.tau.sq %% y[, k] + solve(V, M))
+   }
+   list(pred.mean = pred.mean, pred.var = pred.var,
+     filter.mean = filter.mean, cond.loglik = -0.5 *
+     cond.dev, loglik = -0.5 * dev)
+ }
```

Now we can run it on the example data we generated above.

```
> y <- data.array(ou2)
> a <- matrix(p[c("alpha.1", "alpha.2", "alpha.3",
+   "alpha.4")], 2, 2)
> b <- diag(1, 2)
> sigma <- matrix(c(p["sigma.1"], p["sigma.2"],
+   0, p["sigma.3"]), 2, 2)
> tau <- diag(p["tau"], 2, 2)
> fit2 <- kalman.filter(y, x0, a, b, sigma, tau)
```

In this case, the Kalman filter gives us a log likelihood of `fit2$loglik=-398.535251527343`, while the particle filter gives us `fit1$loglik=-397.766847018944`.

5. THE MIF ALGORITHM

In order to use MIF, we need to specify the distribution of particles in the state-parameter space. This distribution must be such that, when `sd=0`, all the particles are identical. For this example, we'll just draw our particles from a multivariate normal distribution.

```
> normal.particles <- function(Np, center, sd, ivpnames,
+   ...) {
+   params <- matrix(rnorm(Np * length(center),
+     mean = center, sd = sd), length(center),
+     Np, dimnames = list(names(center), NULL))
+   states <- params[ivpnames, , drop = FALSE]
+   rownames(states) <- gsub(".0", "", ivpnames)
+   list(states = states, params = params)
+ }
```

Now we'll run MIF to maximize the likelihood over two of the parameters and the initial conditions. We'll use 1000 particles, an exponential cooling factor of 0.95, and a fixed-lag smoother with lag 10 for the initial conditions.

```
> ou2 <- mif(ou2, Nmif = 0, start = c(x1.0 = 50,
+   x2.0 = -50, p), pars = c("alpha.1", "alpha.4"),
+   ivps = c("x1.0", "x2.0"), particles = normal.particles,
+   rw.sd = c(x1.0 = 5, x2.0 = 5, alpha.1 = 0.1,
+     alpha.2 = 0, alpha.3 = 0, alpha.4 = 0.1,
+     sigma.1 = 0, sigma.2 = 0, sigma.3 = 0,
+     tau = 0), alg.pars = list(Np = 1000, var.factor = 1,
+     ic.lag = 10, cooling.factor = 0.95), max.fail = 100)
```

Just to make it interesting, we'll start far from the true parameter values:

```
> coef(ou2, c("x1.0", "x2.0", "alpha.1", "alpha.4")) <- c(45,
+   -60, 0.8, 0.9)
> tic <- Sys.time()
> fit <- mif(ou2, Nmif = 2, max.fail = 100)
> fit <- continue(fit, Nmif = 78, max.fail = 100)
> toc <- Sys.time()
> print(toc - tic)
```

Time difference of 58.46113 secs

```
> coef(fit)
      x1.0      x2.0      alpha.1      alpha.2      alpha.3
51.2766091 -50.9626991  0.9002916  0.0000000  0.0000000
      alpha.4      sigma.1      sigma.2      sigma.3      tau
 0.9849348  1.0000000  0.0000000  2.0000000  1.0000000
```

The log likelihood of the random-parameter model at the end of the mif iterations, which should be a rough approximation of that of the fixed-parameter model, is `logLik(fit)=-400.419861114764`. To get the log likelihood of the fixed-parameter model (up to Monte Carlo error) we can use `pfilter`:

```
> pfilter(fit)$loglik
[1] -398.8241
```

We can diagnose convergence of the MIF algorithm using “convergence plots” (Fig. 3).

Like `pomp` objects, `mif` objects can be simulated (Fig. 4).

A. A. KING, DEPARTMENTS OF ECOLOGY & EVOLUTIONARY BIOLOGY AND MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MICHIGAN 48109-1048 USA

E-mail address: kingaa at umich dot edu

URL: <http://www.umich.edu/~kingaa>




FIGURE 3. Convergence plots can be used to help diagnose convergence of the MIF algorithm.

```
> plot(simulate(fit)[[1]])
```




FIGURE 4. `mif` objects can be simulated.